OIPE
AUG 28 2002
PATENT & TRADEMARK OFFICE

#35
8/28/02
GJP

PATENT
P54508

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF APPEALS AND INTERFERENCES

In re Application of: HAE-SEUNG LEE **Appeal No.** _____

Serial No.: 08/931,125 Examiner: PORTKA, GARY

Filed: 16 September 1997 Art Unit: 2187
(CPA filed on 27 March 2001)

For: MEMORY SYSTEM FOR IMPROVING DATA INPUT/OUTPUT PERFORMANCE AND METHOD OF CACHING DATA RECOVERY INFORMATION

Attn: Board of Patent Appeals & Interferences

APPEAL BRIEF

RECEIVED
SEP 0 4 2002
Technology Center 2100

Commissioner for Patents
Washington, D.C. 20231

Sir:

Pursuant to appellant's notice of appeal filed on 28 June 2002, appellant (hereinafter "Lee") hereby appeals to the Board of Patent Appeals and Interferences from the final rejection of claims 1 through 8 as set forth in the final office action mailed on 5 March 2002 (Paper No. 30).

Folio: P54508
Date: 8/28/02
I.D.: REB/RHS/kf

08/29/2002 SSITHIB1 00000038 08931125
01 FC:120 320.00 OP

I. REAL PARTY IN INTEREST

The real party in interest is:

SamSung Electronics Co., Ltd.
#416, Maetan-dong, Paldal-gu
Suwon-city, Kyungki-do, Republic of KOREA,

as evidenced by the Assignment executed by the inventor on 2 October 1997 and recorded in the U.S. Patent & Trademark Office on 16 October 1997 at Reel 012849, frame 0001.

II. RELATED APPEALS AND INTERFERENCES

There are no other appeals and interferences known to Appellant, Appellant's legal representatives, or assignee, which will directly affect, be directly affected by or have a bearing on the Board's decision in the pending appeal.

III. STATUS OF CLAIMS

Claims 1 through 8 are pending in this application. All pending claims were finally rejected in the office action mailed on 5 March 2002 (Paper No. 30). Claims 1, 6, and 7 are independent and the remaining claims are dependent.

IV. STATUS OF AMENDMENTS

No Amendment was filed after the final Office action mailed on 5 March 2002

(Paper No. 30).

V. SUMMARY OF INVENTION

As shown in Fig. 1, the RAID-5 system of this invention includes a central processing unit (CPU) 2, a controller 6 connected to the CPU 2 via an input/output bus 4, and a plurality of disk drives DR1-DR5 connected to the controller 6 via SCSI bus 8. CPU 2 transmits data transmitted through an input/output bus 4 from a host computer to controller 6. CPU 2 thereby controls controller 6 to control input/output data between drive disks DR1 to DR5, which are connected to CPU 2 and SCSI bus 8. Each drive DR1 to DR5 connected to SCSI bus 8 records and reproduces the data transmitted from the host computer under the control of controller 6.

Fig. 2 illustrates an example of data transmission of the RAID-5 structure. Data ND transmitted from the host computer is divided by strip (the data is divided by strip 3 in Fig. 2), distributed and stored in each drive DR1 to DR5. That is, each drive DR1 to DR5 has a data block D in which data is stored, and a parity block P in which parity information is stored, to thereby store the data transmitted from the host computer under the control of controller 6.

Fig. 3 is a control flow chart for explaining the writing of the data and parity information transmitted from the host computer in each drive in the RAID structure of

level 5. Referring to Fig. 3, when data writing instruction is received from the host computer, CPU 2 calculates a target location at step 10. At step 12, CPU 2 transmits the data transmitted from the host computer to controller 6. Controller 6 reads old data OD and old parity OP stored in each drive at steps 14 and 16. Next, the controller 6 calculates a new parity NP according to the following formula (1).

$$NP = OP \veebar OD \veebar ND \quad (\veebar \text{ means exclusive OR, i.e., XOR})$$

Controller 6 writes data ND and new parity NP in a predetermined drive at steps 20 and 22. As described, in case a writing instruction of a short data block is received from the host computer in the RAID-5 system, access of another disk on the strip is brought about which contributes to a deterioration to the entire system performance. This appears in the on-line transaction processing environment having many operation loads. That is, in case of the partial strip writing, old parity OP and old data OD are read from a predetermined drive, exclusive-ORed according to formula (1), its result is exclusive-ORed with data ND, and then new parity NP and new data ND are written in a predetermined drive. Thus, two-time reading and writing operations are needed which results in a larger overhead of write data in comparison with a single large expensive drive.

Fig. 4 illustrates a RAID-5 system to which parity cache arrays 38 are connected according to an embodiment of the present invention. This RAID-5 system has a CPU 30 for controlling the overall system. A controller 34 is connected to CPU 30 through an input/output bus 32 to distribute and store data transmitted from a host computer to each drive array 39, or reproduce the stored data under the control of CPU 3. Drives 1 to 5 (39) which are connected to controller 34 through SCSI bus 36 to store and reproduce the data and data recovery information (parity information) transmitted from the host computer under the control of controller 34. Caches 1 to 5 (38) which are connected to controller 34 and input/output bus 36 placed between drives 39 to store the parity information.

Each drive 39 consists of a plurality of blocks in order to store and read the data and parity information. Furthermore, each drive 39 sets up the predetermined number of parity block from the cylinder zero on the disk, and uses it as a parity information storing region, without using the stripping method defined in the RAID structure in level 5. Here, the data cannot be recorded in the parity information storing region.

Fig. 5 is a control flow chart for explaining a process of writing data and parity information in the RAID system constructed according to the embodiment of the present invention. First, the data writing instruction is received from the host computer, CPU 30 updates a task file required at step 40, and then calculates a target cylinder (=parity block

+ request cylinder) in order to use a separate parity block in the drive. Then, CPU 30 transmits new data ND to be written at step 42. Controller 34 next reads old data OD from a predetermined drive 39 in order to generate new parity NP, and then examines if old parity information OP to be read is hit in cache 38 at step 46. Here, if the old parity information OP is hit in cache 38, controller 34 proceeds to step 50. If the old parity information is not hit in cache 38, controller 34 proceeds to step 48. That is, in case that the old parity information OP and parity information are not hit, controller 34 reads the old parity information OP from the predetermined drive at step 48, updates a cache table, and then moves to step 50. Controller 34 calculates a new parity NP by exclusive-ORing the old parity information read and the new data ND through the following formula (2).

$$NP = OP \veebar OD \veebar ND$$

Controller 34 updates the cache table at step 52, and then writes the new data ND transmitted from the host computer and the calculated new parity NP in a predetermined drive at steps 54 and 56. Then, the data writing process of the present invention is completed.

The parity cache is connected between each drive and controller in order to rapidly apply parity information read request. Furthermore, since the parity block for storing the

parity information is set up from the cylinder zero on the disk, it is possible to prevent time delay due to a separate search when sequential read/write operation is carried out.

VI. ISSUES

The ultimate issues are whether claims 1, 2, and 6 through 8 were properly rejected under 35 U.S.C. § 102 for alleged anticipation by Jones U.S. Patent No. 5,572,660; and whether claims 3 through 6 were properly rejected under 35 U.S.C. § 103 for alleged unpatentability over Jones in view of Holland *et al.* U.S. Patent No. 5,455,934. More specifically, however, the issues are as follows:

1. Whether the § 102 and § 103 claim rejections were erroneous because the cited references failed to disclose all elements of the rejected claims.

2. Whether the § 102 rejections of claims 7-8 were erroneous because the Examining Staff failed to apply § 112 ¶ 6 properly in construing the claims.

3. Whether the § 103 rejections of claims 3-5 were erroneous because the Examining Staff failed to document a teaching, suggestion, or motivation in the prior art to combine the cited references.

VII. GROUPING OF THE CLAIMS

The claims do not stand or fall together, and thus the claims are grouped as follows, in accordance with the arguments made below:

Group I: Apparatus claims 1-2 and 6, rejected under § 102.

Group II: Step-plus-function method claims 7-8, rejected under § 102.

Group III: Apparatus claims 3-5, rejected under § 103.

VIII. ARGUMENT

A. THE § 102 REJECTION OF CLAIMS 1-2 AND 6 IS ERRONEOUS BECAUSE THE JONES PATENT DOES NOT DISCLOSE ALL ELEMENTS OF THE CLAIMS

All independent claims in this case are limited to RAID-5 systems having one-to-one caching.[1] Claims 1-2 and 6 stand rejected under § 102 on the ground that the Jones reference teaches one-to-one caching in a RAID-5 system. It is well settled that a § 102 rejection cannot stand if the single reference allegedly supporting the rejection fails to contain each and every element (limitation) of the rejected claim. *In re Robertson*, 169 F.3d 743, 49 USPQ2d 1949 (Fed. Cir. 1999); *In re Schreiber*, 128 F.3d 1473, 44 USPQ2d 1429 (Fed. Cir. 1997); *Gechter v. Davidson*, 116 F.3d 1454, 43 USPQ2d 1030 (Fed. Cir. 1997). If a single element or limitation is missing from the reference, no anticipation occurs. *Motorola, Inc. v. Interdigital Technology Corp.*, 121 F.3d 1461, 43 USPQ2d

[1] Independent claim 1 is directed to "A redundant array of inexpensive disks (RAID) level 5 memory system, comprising: . . . a plurality of caches, each of said plurality of caches respectively coupled operatively to a corresponding single unique one of said plurality of defect-adaptive memory devices, each of said plurality of caches adapted for storing parity information for data recovery for a corresponding single unique one of said plurality of defect-adaptive memory devices to provide one-to-one caching . . . ." Claim 2 depends from claim 1. Claim 6 is an independent apparatus claim but it has substantially the same limitation.

1481 (Fed. Cir. 1997); *Kloster Speedsteel AB v. Crucible, Inc.*, 793 F.2d 1565, 230 USPQ 82 (Fed. Cir. 1986).

The PTO Examining Staff has the burden to show, by a preponderance of evidence, that Lee is not entitled to a patent because the claimed subject matter is anticipated by, or is obvious from, the art of record. Lee is entitled to a patent "unless" the PTO establishes otherwise. See, e.g., *In re Dembiczak*, 175 F.3d 994, 1001, 50 U.S.P.Q.2d 1614 (Fed. Cir. 1999); *In re Epstein*, 32 F.3d 1559, 1564 (Fed. Cir. 1994); *In re Rijckeart*, 9 F.3d 1551, 1552, 24 U.S.P.Q.2d 1443, 1444 (Fed. Cir. 1992); *In re Fine*, 837 F.2d 1071, 1074, 5 U.S.P.Q.2d 1596 (Fed. Cir. 1988).

The Examining Staff made the following arguments in support of the erroneous contention that Jones teaches one-to-one caching in a RAID-5 system:

First, the staff invokes the Abstract of the Jones patent because it says that each data drive is associated with a dedicated write-through cache and it says that the parity drive is associated with a dedicated write-back cache. But the Abstract says *nothing* of RAID-5. Moreover, the assertion that each drive is associated with a dedicated cache does not exclude each drive being associated with more than one dedicated cache. Reciting presence of "an element" never excludes presence of two or three or more elements, in patent drafting. The Abstract proves nothing.

Second, the staff invokes Jones' Fig. 2D. Paper No. 30 states: "Figure 2D clearly

shows one cache for each disk; it would be counter to the teaching of the reference to assume that any cache is connected other than as shown in their one-to one connection." Nothing documentary is cited to support the claim of what the supposed "teaching of the reference" is. Fig. 2D with equal logic is merely illustrative and not exclusive. Also it has ellipsis (...) between the caches and between disk drives. What goes in the ellipsis is left to the imagination of the viewer. Fig. 2D contains nothing that expressly excludes one-to-many connections, such as a cache unit with multiple disk drives or *vice versa*. All that the staff derives from Fig. 2D is mere speculation. Fig. 2D does not expressly speak to the issue in controversy here. Paper No. 30's reliance on Fig. 2D therefore does not satisfy the requirement of substantial evidence, per *In re Zurko*, 258 F.3d 1379, 1385-86 (Fed. Cir. 2001) ("With respect to core factual findings in a determination of patentability" PTO "must point to concrete evidence in the record" and "cannot simply reach conclusions based on its own understanding or experience").

What is even more to the point than Fig. 2D is the text of the specification of Jones. That is where one would look for Jones to state what he proposes to teach. Jones discusses Fig. 2D and RAID-5 in column 10, starting at line 7. Significantly, Jones says nothing about maintaining a one-to-one caching scheme for RAID-5 in the implementation described. Thus, Jones says:

> ...[T]he parity information is stored and distributed among the plurality of disk drives 214-1 through 214-8 according to a level 5 RAID approach. As such, a

> plurality of cache units 254-1 through 254-8 are coupled between array scheduler 210 and disk drives 214-1 through 214-8. The cache units 254-1 through 254-8 are partitioned and configured such that a write-through caching scheme is carried out when array scheduler 210 writes real data to a selected disk drive and such that a write-back caching scheme is carried out when array scheduler 210 writes parity information to a selected disk drive.

There is not a word about one-to-one caching in that passage.

Lee respectfully submits that nothing but conjecture supports the claim by the Examining Staff that Jones teaches one-to-one caching in a RAID-5 system. Accordingly, Lee respectfully relies on the command of the Federal Circuit's recent decisions in *In re Lee*, 277 F.3d 1338 (Fed. Cir. 2002), and *In re Zurko*, 258 F.3d 1379 (Fed. Cir. 2001). Those decisions insist on specificity and documentation. They do not permit speculation and conjecture. They do not permit so-called common knowledge or the introspection of the Examining Staff to be the basis of rejections. They demand substantial evidence on the record. Thus *Zurko*, 258 F.3d at 1385-86 states: "With respect to core factual findings in a determination of patentability" the PTO "must point to concrete evidence in the record" and "cannot simply reach conclusions based on its own understanding or experience—or on its assessment of what would be basic knowledge or common sense." Under *Lee* and *Zurko*, the rejection here is unsupported by substantial evidence and therefore should be reversed.

The Jones patent does not disclose all elements of the rejected claims, namely one-to-one caching in a RAID-5 system. Therefore, the Examining Staff has not met its

burden under § 102.

B. THE § 102 REJECTION OF CLAIMS 7-8 IS ERRONEOUS BECAUSE THE EXAMINING STAFF FAILED TO APPLY § 112 ¶ 6 PROPERLY

Independent claim 7 (from which claim 8 depends) is directed to a method "of writing data to, and reading data from, a redundant array of inexpensive disks (RAID) level 5 system." The claim is in Jepson format and it ends with the words -- "*the improvement comprising* a step for reducing overhead during a read operation for data recovery and thereby improving data input-output performance." The language expressly uses the phrase "a step for" and the remaining claim language is entirely functional -- reducing overhead and improving data input-output performance. This language invokes § 112 ¶ 6.

Therefore, the specification is a glossary for the claim. *Chiuminatta Concrete Concepts v. Cardinal Indus., Inc.*, 145 F.3d 1303, 46 U.S.P.Q.2d 1752 (Fed. Cir. 1997). However, the Examining Staff asserted (p. 6, ¶ 14):

> Applicants have argued that Jones does not teach a unique one cache corresponds to a unique one disk, in a one-to-one caching for a RAID 5 system. ... It is further noted that this argument is not supported by the language of claim 7.

The argument of Lee, which is the same as that which Lee made above that the rejection fails under the all-elements rule, is fully supported by the language of independent claim 7. As a step-plus-function claim, claim 7 incorporates (by virtue of § 112 ¶ 6) the acts or structure in the specification that performs the recited functions.

Those acts or structure clearly include the use of one-to-one caching in the "writing [of] data to, and reading [of] data from, a redundant array of inexpensive disks (RAID) level 5 system."

As shown above, the cited Jones patent does not teach one-to-one caching for a RAID-5 system. Therefore the reference does not have all elements (limitations) of the claim, and the § 102 rejection fails.

Paper No. 30 also states (p. 4, ¶ 7) that Jones anticipates Lee because Jones "teaches a step for reducing overhead during a read for data recovery by avoiding the need to access the disk two times when the required data is in a cache." Although language similar to that just quoted constitutes the recited function of claim 7, that is not the proper legal test under § 112 ¶ 6.

It is not enough for § 102 purposes that the cited reference performs the same *function* as the claim recites. The Federal Circuit has held it error to assume that two structures are the same or equivalent simply because they perform the same function. *Roton Barrier, Inc. v. Stanley Works,* 79 F.3d 1112, 1126-27 (Fed. Cir. 1996); *Pennwalt Corp. v. Durand-Wayland, Inc.,* 833 F.2d 931, 934 (Fed. Cir. 1987) (en banc) ("Pennwalt erroneously argues that, if an accused structure performs the function required by the claim, it is per se structurally equivalent"), *cert. denied,* 485 U.S. 961 (1988). Infringement (or anticipation) is found only if the claimed function is performed by either the

same structure (or acts) that the specification describes or else by an equivalent of the structure (or acts). *Texas Instruments Inc. v. United States Int'l Trade Comm'n*, 805 F.2d 1558, 1562, 231 USPQ 833, 834-35 (Fed. Cir. 1986).

The same principle applies to method step-plus-function claims as it does in apparatus means-plus-function claims. Under § 112 ¶ 6, the necessary sameness or equivalency must be supported by substantial evidence of record. This record does not contain any evidence of such sameness or equivalency. Therefore, the § 102 rejection of these claims fails.

It does not matter whether Jones teaches a step for accomplishing the recited function. In this regard, Lee respectfully points out also that it would not be enough for anticipation even if Jones reduced overhead "by avoiding the need to access the disk two times when the required data is in a cache." For anticipation, Jones would also have to perform this function in substantially the same way as Lee does or do so in a manner insubstantially different from how Lee does. Jones does not perform these acts in substantially the same way as Lee does, or in just an insubstantially different way, and therefore Jones does not anticipate Lee.

C. THE § 103 REJECTION OF CLAIMS 3-5 IS ERRONEOUS BECAUSE THE COMBINED REFERENCES DO NOT DISCLOSE ALL ELEMENTS OF THE CLAIMED SUBJECT MATTER

Claims 3-5, dependent from claim 1 directly or indirectly, were rejected under 35

U.S.C. § 103(a) for alleged obviousness over Jones in view of Holland *et al.* U.S. Patent No. 5,455,934. As discussed above regarding claim 1, Jones does not teach use of one-to-one disk caching in a RAID-5 system, and such remarks are incorporated herein by reference. Holland does not supply the limitation.

All elements must be present in a combination of references relied upon under § 103 just as they must be in the single reference relied upon under § 102. *In re Lowry*, 32 F.3d 1579, 32 U.S.P.Q.2d 1031 (Fed. Cir. 1994). The lack of all claim elements, in Jones and Holland collectively, undermines the obviousness rejection of claims 3-5.

D. The § 103 Rejection of Claims 3-5 Is Unsupported Because No Teaching, Suggestion, or Motivation in the Prior Art Was Established for Combining the Two References

Claims 3-5, dependent from claim 1 directly or indirectly, were rejected under 35 U.S.C. § 103(a) for alleged obviousness over Jones in view of Holland *et al.* U.S. Patent No. 5,455,934. The Examining Staff concedes (Paper No. 30, ¶ 11) that claim 3 requires, and Jones does not disclose, "that the information needed for data recovery is sequentially arranged from the most outer cylinder." (Claims 4 and 5 depend from claim 3 and therefore they have the same limitation.)

It is well settled that, before the PTO may combine the disclosures of two or more prior art references in order to establish *prima facie* obviousness, it must establish on the record that some specific suggestion, motivation, or teaching is found in the prior art that

would have led an ordinary artisan to select those specific references and to adapt and combine them in the same way that the inventor did. *Karsten Mfg. Corp. v. Cleveland Gulf Corp.*, 243 F.3d 1376, 1385, 58 USPQ2d 1286, 1293 (Fed. Cir. 2001) ("In holding an invention obvious in view of a combination of references, there must be some suggestion, motivation, or teaching in the prior art that would have led a person of ordinary skill in the art to select the references and combine them in the way that would produce the claimed invention."); *In re Dembiczak*, 175 F.3d 994, 999, 50 U.S.P.Q.2d 1614 (Fed. Cir. 1998) (teaching or motivation or suggestion to combine is an "essential evidentiary component of an obviousness holding"); *ATD Corp. v. Lydall, Inc.*, 159 F.3d 534, 546, 48 U.S.P.Q.2d 1321 (Fed. Cir. 1998)("There must be a teaching or suggestion within the prior art, or within the general knowledge of a person of ordinary skill in the field of the invention, to look to particular sources of information, to select particular elements, and to combine them in the way they were combined by the inventor."); *In re Rouffet*, 149 F.3d 1350, 1355, 47 U.S.P.Q.2d 1453 (Fed. Cir. 1998); *In re Chu*, 66 F.3d 292 (Fed. Cir. 1995).

The Examining Staff tried to supply the necessary suggestion, motivation, or teaching in the prior art to combine with Jones' disclosure the element "that the information needed for data recovery is sequentially arranged from the most outer cylinder" by reliance on what was allegedly well known in the art and teachings from the

Jones and Holland patents. Paper No. 30, ¶ 11, states:

> However, it is well known that the sequential nature of disk access invites a transfer mechanism sequentially from some position, thus improving performance by reducing seek time. As further taught by Holland, arrangement of information on a disk from the outermost cylinders results in higher sustained data transfer rates (see column 9 lines 25-30). It is clear from Jones at column 2 lines 34-58 that the accessing of the parity data in RAID systems limits the performance of these systems, and therefore the advantage of faster access due to reduced seek time, and higher sustained data rates would have motivated an artisan to arrange this information from the outermost cylinder. Thus it would have been obvious to one of ordinary skill in the art at the time of the invention to sequentially arrange the recovery information from the most outer cylinder in Jones, because this method reduces seek time, results in higher sustained data rates, and therefore improves performance.

First, Lee has challenged the undocumented assertion that "it is well known" that the nature of data access "invites" the described sequential transfer. Lee demanded a citation to a text or other authority to support the claim of what is well known (see Paper No. 29, sec. II), but the Examining Staff has supplied none. This failure to document an essential premise of the rejection requires reversal. See *In re Lee*, 277 F.3d 1338 (Fed. Cir. 2002), holding that core underpinnings of the basis for combining references must be documented from the prior art, *not* left for speculation or conjecture, and *not* "resolved on subjective belief and unknown authority." *Id.* at 1344. See also *In re Zurko*, 258 F.3d 1379, 1385-86 (Fed. Cir. 2001) ("With respect to core factual findings in a determination of patentability" PTO "must point to concrete evidence in the record" and "cannot simply reach conclusions based on its own understanding or experience—or on its assessment of what would be basic knowledge or common sense").

Second, the combination of the allegedly well-known invitation from "the nature of data access" with the teachings of Jones and Holland is not supported by any specific teaching, suggestion, or motivation in the prior art to make the combination. The required teaching, suggestion, or motivation in the prior art is not supplied by the conclusion of Paper No. 30 that making the combination "reduces seek time, results in higher sustained data rates, and therefore improves performance." That desirable result may occur, but in that event the desirable result is only what you will get if you combine the references, *not* what will suggest, motivate, or teach you (if you are without knowledge of Lee's invention) to combine the references. This is a classic example of the 20-20 hindsight and bootstrapping that such decisions as *Lee* and *Dembiczak* command the Examining Staff to avoid. You have to supply the justification for combining the selected elements of the references *before* you combine them, in order to develop the required motivation or suggestion; you cannot operate in the reverse direction, by inferring motivation to combine from the benefits accomplished by the combination of elements once it is made.

IX. CONCLUSION

In view of the law and facts stated herein as well as all the foregoing reasons, Appellant believes that the rejection is improper and respectfully requests that the Board refuse to sustain the outstanding rejection of claims 1 through 8 under 35 U.S.C. §§ 102 and 103.

Respectfully submitted,

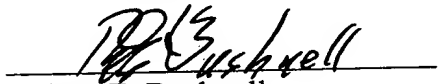

Robert E. Bushnell,
Attorney for the Applicant
Registration No.: 27,774

1522 K Street N.W., Suite 300
Washington, D.C. 20005
(202) 408-9040

Folio: P54508
Date: 8/28/20
I.D.: REB/kf

X. APPENDIX

CLAIMS UNDER APPEAL (1-8)

1 (five times amended). A redundant array of inexpensive disks (RAID) level 5 memory system, comprising:

a plurality of defect-adaptive memory devices, each of said plurality of defect-adaptive memory devices having a first region for sequentially storing parity information for data recovery and a second region for storing data;

a plurality of caches, each of said plurality of caches respectively coupled operatively to a corresponding single unique one of said plurality of defect-adaptive memory devices, each of said plurality of caches adapted for storing parity information for data recovery for a corresponding single unique one of said plurality of defect-adaptive memory devices to provide one-to-one caching; and

a controller operatively coupled to each defect-adaptive memory device of said plurality of defect-adaptive memory devices and to each corresponding single unique cache of said plurality of caches, said controller comprising a first means for selectively controlling writing and reading of parity information needed for data recovery in said first region of each corresponding single unique one of said plurality of defect-adaptive memory devices, a second means for selectively obtaining parity information needed for data recovery from said first region of each corresponding single unique one of said plurality of defect-adaptive memory devices, and a third means for selectively storing parity information needed for data recovery obtained from said first region of a corresponding single unique one of said plurality of defect-adaptive memory devices in a predetermined corresponding single unique one of said plurality of caches.

2 (twice amended). The memory system of claim 1, wherein said controller comprises a means for determining whether data recovery information is stored in any cache of said plurality of caches.

3 (thrice amended). The memory system of claim 1, wherein the parity information needed for data recovery is stored and is sequentially arranged from the most outer cylinder on a recording medium in each corresponding one of said plurality of defect-adaptive memory devices.

4 (thrice amended). The memory system of claim 3, wherein parity information for data recovery is modified to a value obtained through a calculation of new data recovery information.

5 (thrice amended). The memory system of claim 4, wherein parity information for data recovery is obtained by performing an exclusive-OR operation on previous data, parity information corresponding to the previous data, and new data.

6 (five times amended). A redundant array of inexpensive disks (RAID) level 5 system, comprising:

a plurality of disk drives, each of said plurality of disk drives including a first region having a plurality of data blocks for storing data and a second region having a predetermined number of parity blocks for storing parity information for data recovery;

a plurality of caches, each of said plurality of caches respectively coupled operatively to a corresponding single unique one of said plurality of disk drives, each of said caches

adapted for storing parity information for data recovery; and

a controller adapted to provide one-to-one caching, said controller operatively coupled to each disk drive of said plurality of disk drives and to each corresponding single unique cache of said plurality of caches, said controller adapted for selectively controlling a write operation of data and parity information for a data recovery in each corresponding disk drive of said plurality of disk drives, said controller comprising:

a first means for selecting a single predetermined disk drive of said plurality of disk drives upon receipt of a data writing instruction from a host computer;

a second means for reading old data from the single predetermined disk drive of said plurality of disk drives;

a third means for determining whether old parity information corresponding to the old data corresponding to the single predetermined disk drive of said plurality of disk drives is accessed in a corresponding single unique cache of said plurality of caches;

a fourth means for reading the old parity information from the single predetermined disk drive of said plurality of disk drives, upon the old parity information corresponding to the single predetermined disk drive of said plurality of disk drives not being accessed in the corresponding single unique cache of said plurality of caches, and for then loading the corresponding single unique cache of said plurality of caches with the old parity information;

a fifth means for obtaining new parity information by performing an exclusive OR operation on the old data, the old parity information and

new data;

a sixth means for loading the corresponding single unique cache of said plurality of caches with the new parity information; and

a seventh means for writing the new data in said region for storing data in the single predetermined disk drive of said plurality of disk drives and writing the new parity information in said another region for storing parity information in the predetermined single disk drive of said plurality of disk drives,

whereby the data writing process is completed.

7 (amended). In a method of writing data to, and reading data from, a redundant array of inexpensive disks (RAID) level 5 system, said method comprising steps for sequentially storing information for data recovery in a first region of a disk, storing information comprising data in a second region of the disk other than the first region, controlling writing and reading of information by means of an electronic controller unit, and caching information for data recovery, *the improvement comprising* a step for reducing overhead during a read operation for data recovery and thereby improving data input-output performance.

8 (twice amended). The method of claim 7, wherein said step for reducing overhead during a read operation for data recovery and thereby improving data input-output performance comprises steps for:

(a) coupling each one of a plurality of caches to each corresponding one of a plurality of disks, whereby each disk is coupled one-to-one to one cache;

(b) operatively coupling the caches to the controller;

(c) storing in each one of the plurality of caches information for data recovery in the disk corresponding to the cache; and

(d) determining information for data recovery in a disk by using information for data recovery stored in the cache corresponding to the disk.